AIAA 2003–3698

Aircraft Gas Turbine Engine Simulations

W. C. Reynolds , J. J. Alonso, and M. Fatica

Center for Integrated Turbulence Simulations
Stanford University, Stanford, CA 94305

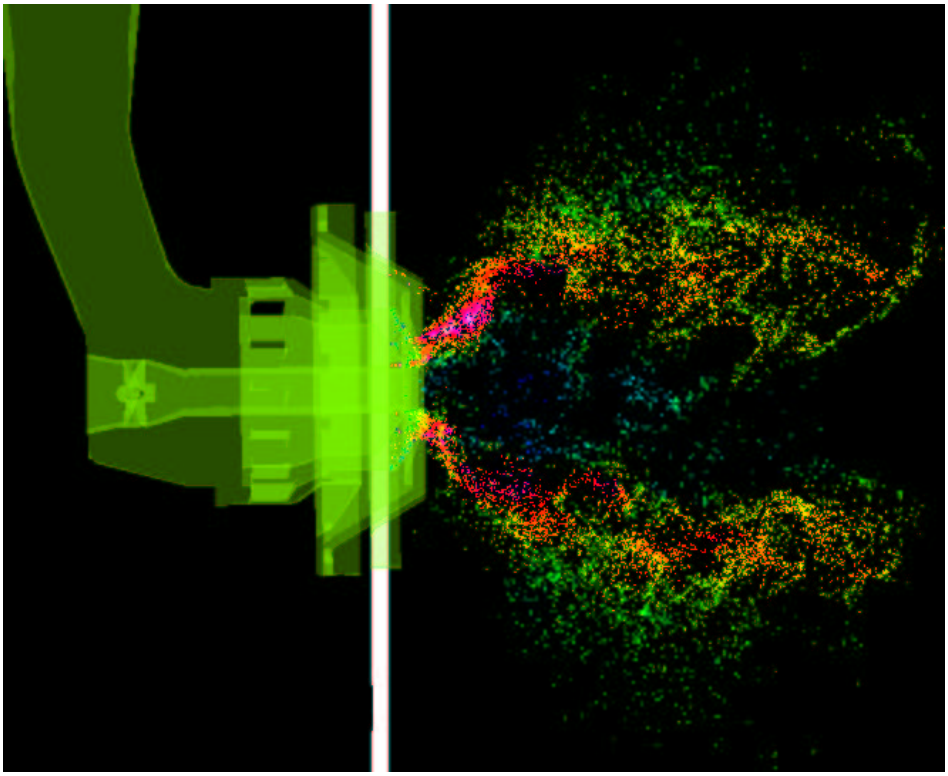

16th AIAA Computational Fluid Dynamics Conference
June 23–26, 2003/Orlando, FL

AIAA 2003–3698

Aircraft Gas Turbine Engine Simulations

W. C. Reynolds*, J. J. Alonso†, and M. Fatica‡

Center for Integrated Turbulence Simulations
Stanford University, Stanford, CA 94305

The Center for Integrated Turbulence Simulations (CITS) was created in October of 1997 under the sponsorship of the Department of Energy (DoE) Accelerated Strategic Computing Initiative (ASCI) Alliance program. The over-arching problem to be solved by this Center is the detailed, high-fidelity simulation of the flowpath through complete aircraft gas turbines including the compressor, combustor, turbine, and secondary flow systems. In order to support this goal, three groups were created to tackle the main components of the program. The turbomachinery group has developed a multiblock-structured, unsteady Reynolds-Averaged Navier-Stokes (URANS) solver called TFLO which is used for the rotating components of the flowpath (compressor and turbine) and for the secondary flow system. The combustion group is responsible for the flow in the combustor and compressor diffuser. Given the complexity of the flow in this region of the engine, an unstructured Large Eddy Simulation (LES) solver called CDP has been developed and validated together with various approaches for combustion simulation and spray tracking. A significant portion of the work in the Center is the integration between the various physics and components being modeled. For this purpose, an integration group is working towards the development of interface definitions for arbitrary URANS and LES solvers. An important aspect of the problems to be solved is that they are of such magnitude and complexity that they require the largest available supercomputers. For this reason, a major thrust of the Center is on the parallel implementation of the solver software and its integration on massively parallel computing platforms. This paper presents some of the major results from each of the three groups in the Center.

Introduction

THE goal of the DoE's ASCI Alliance program is to push the development of comprehensive numerical simulation tools for important engineering problems of interest to the DoE Defense Programs. The work at the five Alliance universities is associated with physical modeling and software development for a particular unclassified application chosen by that university. The Stanford application is the flow and combustion in aircraft gas turbine engines.

Fig. 1 shows the Pratt & Whitney 6000 engine, an advanced modern engine that provides the focus for our simulations. The large axial flow fan at the front of the engine will be familiar to anyone who has flown a modern commercial jet aircraft. It moves most of the air that provides the engine thrust. The fan is driven by a gas turbine engine, consisting of a two-stage compressor that raises the air pressure, a combustor that heats the air, and a two-stage turbine that extracts mechanical energy from the air. The first high-pressure high-speed turbine drives the compressor, and the subsequent low-pressure low-speed turbine drives the fan. The flow exiting the turbine at relatively high velocity provides additional thrust to augment that produced by the fan. In a modern engine like this, a significant fraction of the compressor flow bypasses the combustor and flows inside the high-pressure turbine blades to cool them, and then out through the blade surface to provide a cool protective film that prevents melting of the blades by the very hot combustor exhaust. All of this presents a very challenging engineering problem for which accurate simulations of the flow and combustion and their interactions with the structure are very important. The primary goal of the simulation component of the Stanford ASCI program is to make a major advance in this simulation technology.

Turbulence is important in all components of an aircraft engine. The Stanford ASCI program is closely coupled to fundamental research in turbulence physics, modeling, and simulation being conducted in the Stanford/NASA Center for Turbulence Research (CTR). The coupled CTR/CITS efforts on turbulence simulation are the most substantial set of such activities in the USA and are among the very best such activities worldwide. The group is pioneering major developments in turbulent combustion physics and simulation, and is also contributing advanced turbulence models for non-reacting turbulent flow analysis. This will be evident in the brief outline of the scope of the engine simulations that follows and in the details presented in the subsequent sections.

*Professor Emeritus, AIAA Associate Fellow
†Assistant Professor, AIAA Member
‡Senior Research Engineer

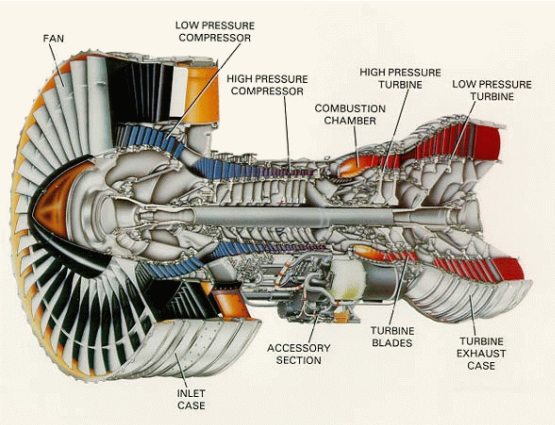

Fig. 1 The PW6000 aircraft gas turbine engine.

Engine Component Simulation

The engine simulation involves two main efforts now being integrated as we move forward to full engine simulation. The turbomachinery is modeled using the unsteady Reynolds-Averaged Navier Stokes (RANS) equations that employ advanced turbulence models being developed at Stanford as part of this and other sponsored research. The turbomachinery code TFLO is a derivative of unsteady flow codes previously developed. The combustor is modeled using Large Eddy Simulation (LES), which has been developed with the conservation properties essential for LES.

TFLO has proven to be comparable in performance and physical content with the best industrial turbomachinery codes, and with coming improvements in algorithms TFLO is expected to become superior to the industrial codes over the next few years.

RANS combustor codes do not do a good job of modeling turbulent mixing in the very complex flow typical of combustor chambers, and our LES code (here designated as CDP) has already proven to be much more accurate in this regard. The LES combustor code is rapidly taking combustor simulation technology to a very high new level. Others have written LES codes for combustion, but these codes lack the conservative, non-dissipative property of our codes that is essential for accurate turbulent flow prediction. Considerable interest in this code is now evident on the part of the aircraft engine industry, particularly Pratt & Whitney, which is collaborating closely with our team on applications to their combustor geometries.

Engine Code Integration

The RANS turbomachinery and LES combustor codes are being integrated to enable coupled simulation of the turbomachinery and combustor. As the available computer power increases, it should become possible to do a very detailed integrated simulation of an entire aircraft engine. This capability would revolutionize the design and development processes and associated cost for aircraft gas turbine engines and many other systems.

The turbomachinery code TFLO was developed rapidly as a derivative of existing codes. However, the development of the unstructured mesh LES combustor code started from scratch with this program, and it has taken more time to bring it to the point where it can be integrated with TFLO. In order to get some experience with code integration, we have integrated TFLO with a RANS combustor code, the National Combustor Code (NCC) developed at NASA/Glenn Research Center. In collaboration with researchers from NASA/Glenn, we have developed an interface that can couple TFLO and the NCC. This framework is now used (almost unaltered) to integrate TFLO and CDP and therefore, valuable integration experience was gained with the all-RANS integrated simulations.

Turbomachinery Simulations

The objective of the turbomachinery group is to simulate complex unsteady flows in the rotating components of a jet engine both accurately and efficiently. The flows around the rotating blades of compressors and turbines are characterized by a number of unsteady physical phenomena that require the solution of the time-dependent Reynolds-Averaged Navier-Stokes equations with a suitable turbulence model. Turbulent boundary layers and secondary flows, transition, flow separation, heat transfer, blade-wake and blade-shock interactions, and other effects, dictate the need for nonlinear fluid dynamics and turbulence models that can resolve fluctuations with disparate temporal and spatial scales. For these reasons, the accurate computation of these unsteady flows requires long time integrations on very large meshes with substantial resolution.

In order to simulate these flows, the turbomachinery group has developed a solver, TFLO, (Turbomachinery FLO solver), that is capable of achieving both high single-processor performance and scalable parallel efficiency using a FORTRAN 90/95, MPI-based implementation. A number of ingredients contribute to making TFLO a state-of-the-art solver. In order to reach a level of maturity that is now comparable with industrial solvers, we have developed an efficient baseline algorithm that integrates an explicit, multigrid, Runge-Kutta solver for the mean flow with both ADI (Alternating Direction Implicit) and DDADI (Diagonally Dominant ADI) solution methodologies for the turbulence model equations. Time integration is carried out using an implicit dual-time stepping scheme, and a sliding mesh interface with conservative interpolation is used to treat the problem of relative blade motion.

Despite great improvements in the performance of turbulence model implementations, parallel efficiency and parallel I/O, and the dual-time stepping time-

integration, and despite the fact that the efficiency of TFLO can be considered to be state-of-the-art, typical unsteady computations of the flow in complete compressors and turbines on current supercomputers still require several months of runtime using large numbers of processors. In order to be able to use tools such as TFLO not only for validation/verification but also for new scientific discoveries and alternative engine component designs, the turnaround time of these calculations must be reduced by over an order of magnitude through algorithmic improvements alone. The increasing processor speeds, together with computers with larger numbers of processors, will contribute the remaining factor needed to complete these calculations with overnight turnaround.

For these reasons, a large portion of our work during the last year has focused on the study of a number of algorithmic alternatives that can improve the steady-state and unsteady performance of TFLO by an order of magnitude. Our progress with four candidate algorithms has been reported previously and can be found in the Center Technical Report.[2] Only a brief description of this work is presented here.

In addition, we have continued to improve the quality of the predictions from TFLO by concentrating on a number of basic science issues which will also be discussed in this paper briefly. Finally, the results of the use of TFLO in some key large-scale simulations of flows in turbines are shown, together with the physical insight gained from the post-processing of these calculations.

Large-Scale Turbomachinery Simulations

The process of developing the capability to perform integrated, multi-component simulations of major or whole sections of a gas-turbine engine includes intermediate steps. These steps consist of validation, verification, and demonstration of the ability to compute the flow through major sections or entire components of the engine, such as the complete high- or low-pressure turbine or combustor. This section focuses on one of the many such efforts that have been tackled by the turbomachinery group during the past year.

TFLO Validation - Transonic Turbine Rig

Validation of the TFLO code for multi-stage, unsteady flow has continued primarily for turbines. During the past year, the validation of TFLO for a single-stage, transonic, high-pressure turbine, transition duct, and 1st-vane of the low-pressure turbine configuration was completed and the results are discussed in this section. Details can be found in the work of Davis.[3] Detailed comparisons between the predicted and experimental fundamental and 1st harmonics of the unsteady airfoil surface pressure spectrum were performed. One quarter of the rig circumference using 9-HPT-vanes, 14-HPT-blades, 9-LPT-vanes was modelled in the investigation. The configuration was discretized with over 32 million grid points and the unsteady-flow solution was obtained with 192 processors of the DoE Blue Pacific IBM SP2 at Lawrence Livermore National Laboratory. This investigation showed that the TFLO code predicted the unsteady surface pressure non-dimensionalized by the upstream total pressure to within 0.6 percent. Comparisons between the TFLO results with the predictions from a similar Pratt & Whitney in-house code showed excellent agreement.

Fig. 2 illustrates the various phenomena that lead to the unsteady interactions between the three components, especially those between the blade, transition duct, and low-pressure turbine vane. These figures show the instantaneous entropy and pressure contours at the 10% span, mid-span, and 90% span planes. The pressure contours shown in Fig. 2 illustrate the pressure jump associated with the blade trailing edge shock system. As the blade rotates, the blade trailing edge shock extends to the downstream vane where it strikes and sweeps forward along the vane pressure side. The blade tip radius is approximately 31.6 cm. The shock sheet generated from the blade trailing edge shock system propagates downstream at a constant radius so that only radial positions below 31.6 cm are directly struck by this shock. The peak unsteady pressure on the vane pressure side is located near the leading edge and around 31.6 cm for this reason. Some moderate levels of unsteady pressure are also located on the vane near and above mid-span. The source of this unsteady pressure is believed to be from blade trailing edge shock reflections off of the transition duct into the vane. Another area of high unsteady pressure is at the vane tip leading edge near the 90% span location. The unsteady pressures at this location can be attributed to the blade tip vortex that migrates inboard somewhat as it convects through the transition duct.

As shown in Fig. 2, the wakes at mid-span from the upstream first vane mix with the blade wakes into bands of high entropy that are clocked to the leading edge of the downstream low-pressure turbine vane. Several recent investigations into blade wake clocking in multi-stage turbomachinery suggest that the current vane wake clocking could provide optimal performance. The blade wakes strike the downstream vane at the leading edge and, along with the first vane wakes, tend to migrate to the suction side of the low-pressure turbine vane. At the 10% and 90% span locations, the secondary and tip vortex flows mix with the first vane wakes more thoroughly such that the entropy going into the second vane is more uniform. The blade tip vortex strikes the low-pressure turbine vane leading edge and migrates to the pressure side of the vane. There are typically two blade/tip vortex wakes

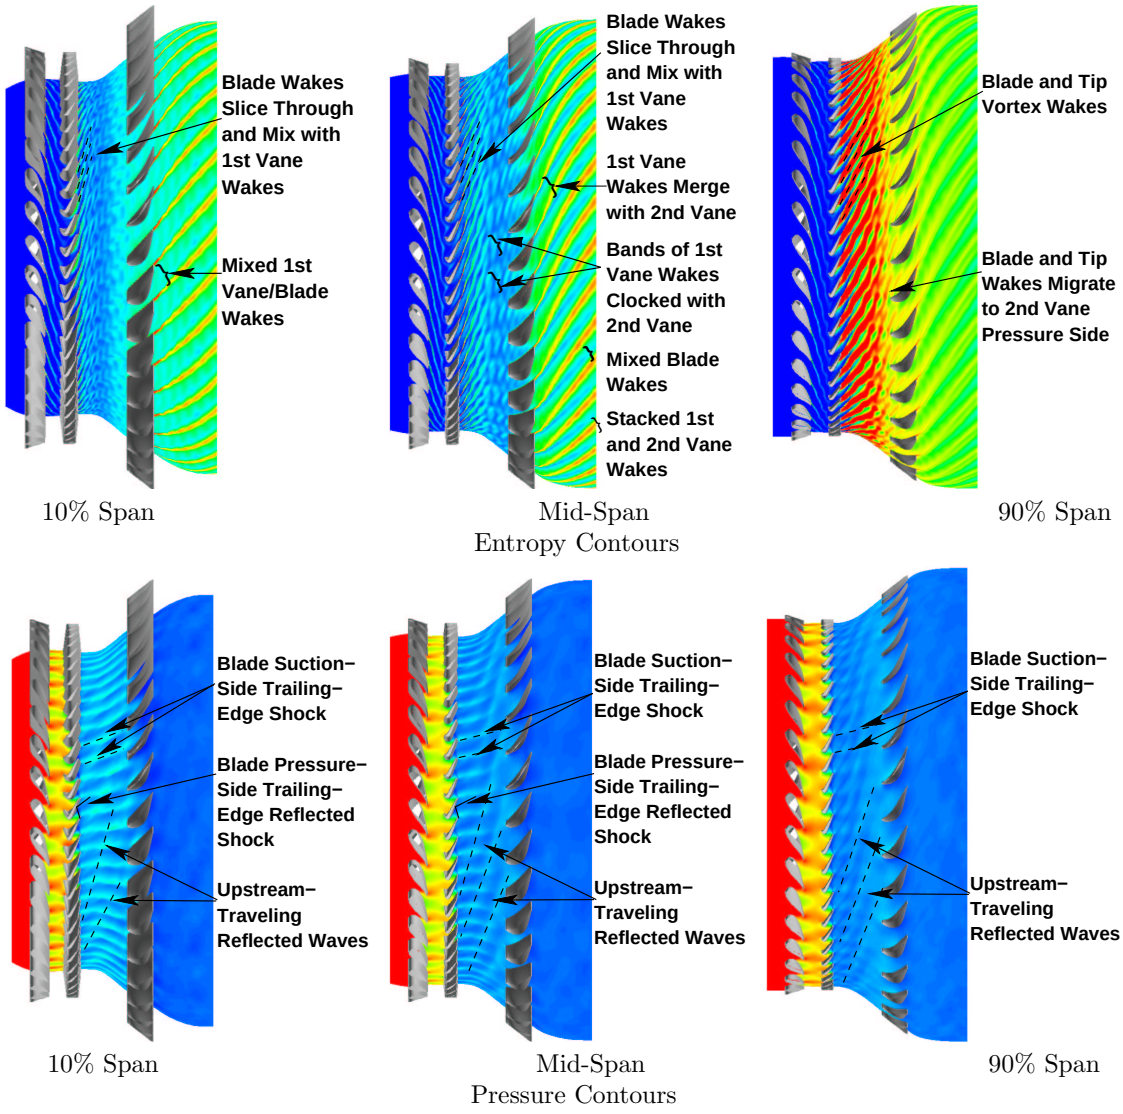

Fig. 2 Computed Instantaneous Entropy and Pressure Contours at 10%, 50% and 90% Span for a Transonic HPT, Transition Duct, 1st-Vane LPT Rig Configuration

in the low-pressure turbine vane passage at any instant in time.

The unsteady pressure signals on the blade surfaces were post-processed using Fast Fourier Transforms (FFT) to obtain information in the frequency domain. Clark[4] showed that the fundamental and first harmonic of the pressure make up over 90% of the pressure envelope over most of the surface.

A comparison between the computed and experimental unsteady surface spectrum at points on the blade surface where the amplitude of the unsteady pressure was also performed. The computed pressure spectra are in good agreement with the experimental data. On the pressure side of the blade, the maximum differences between the computed and experimental amplitudes of the fundamental and first harmonic are around 0.1% of the inlet total pressure. The amplitude of the fundamental mode is slightly under-predicted whereas the amplitude of the first harmonic is slightly over-predicted.

The agreement between the computed and experimental pressure spectrum at the 30% axial chord position is fairly good with the maximum difference being around 0.25% of the inlet total pressure. Once again, the amplitude of the fundamental mode is under-predicted whereas the amplitude of the first harmonic is over-predicted. Larger differences exist between the computed and experimental amplitudes at the 75% axial chord position, however. Here, the simulation over-predicts the fundamental mode amplitude by 0.2% of the inlet total pressure. The difference in the amplitude of the first harmonic is larger (0.6% of the inlet total pressure).

Transition Simulation and Modeling

Transition from laminar to turbulent flow is an important phenomena that has been approached traditionally by instability theory. When transition pro-

ceeds from inflectional shear flow velocity profiles, this is a sound conception: the breakdown to turbulence is the endpoint that evolves from initial instability waves. However, in boundary layers that scenario seems uncommon. Tollmien-Schlichting waves are rarely the precursors of transition, except in carefully controlled wind tunnel experiments. Tollmien-Schlichting waves have very small growth rates and are usually preempted by other mechanisms. Barring inflectional, or other inviscid instability, transition in internal flows often is caused by external disturbances. This scenario is called bypass transition. A portion of our effort within the ASCI program has been to study bypass transition by Direct Numerical Simulation (DNS), and to develop models that can be used in TFLO to predict transition in turbomachinery. In addition, a number of other investigations into the basic science issues in the solution of the URANS equations have been discussed in Ref.[2]

Turbine Passage

We will present some results for three inlet conditions. The first is turbulence free. In that case a simulation with incident wakes was restarted with no inlet disturbance; hence, it started with a disturbance in the free-stream, that washed out, leaving behind a flow with no free-stream turbulence, but with a self-sustained transition near the trailing edge.

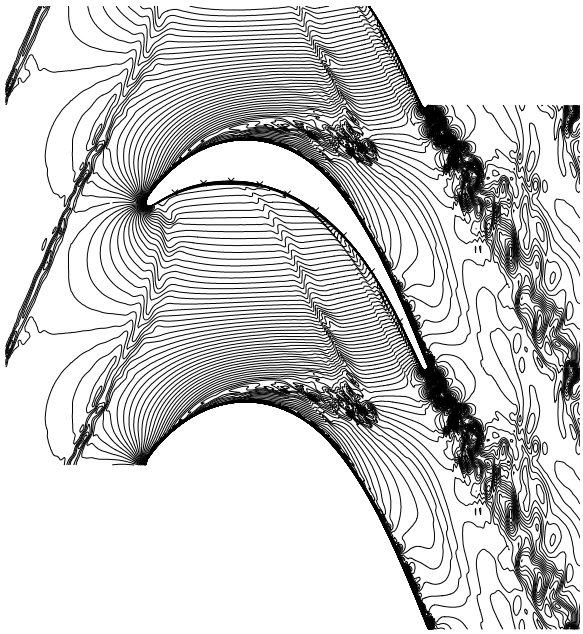

Fig. 3 Magnitude of velocity for case of incident wakes.[5]

The second inflow consists of periodically passing wakes, as shown in Fig. 3. These consist of fully-developed, turbulent wakes obtained from a separate simulation.

The third inflow condition is a field of grid turbulence carried with the flow.

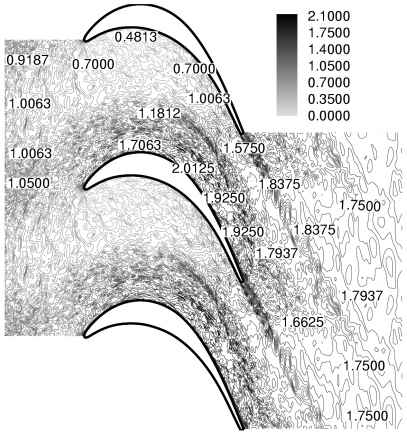

Fig. 4 Magnitude of velocity for case of grid turbulence

The influence of the different inlet conditions on the state of the boundary layer is illustrated by Fig. 5. In the first case transition occurs near the trailing edge, in a region of adverse pressure gradient. The chord Reynolds number is 1.5×10^5 in all cases — which is typical of the low pressure turbine in a full scale engine. This is sufficiently low that most of the blade would be laminar were it not for the inlet disturbance. In this case instability waves are observed near the trailing edge, suggesting that an orderly, Kelvin-Helmholtz instability leads to transition.

With incident wakes a buffeted region is seen near the leading edge. The boundary layer relaminarizes further along the surface; indeed, it is likely that the turbulence seen in this region is driven by the incident wake, rather than being boundary layer turbulence. Shortly past midspan boundary layer transition begins.

Bypass transition occurs in the lowermost case as well. The incident grid turbulence seems to produce a straighter transition line. At this stage it is unclear what causes the qualitative difference between the second and third cases. Clearly the intermittent wake lets the turbulent zone convect toward the trailing edge before being reinvigorated by the subsequent wake. This might be why the transition region has a different appearance to the blade continuously perturbed by grid turbulence.

TFLO Algorithm Improvements

Since the default time-integration scheme in TFLO is based on an implicit, dual-time stepping algorithm which solves a steady-state problem for each physical time step, improvements in the unsteady performance of TFLO can be obtained in one of two ways. One can either improve the performance of the inner iter-

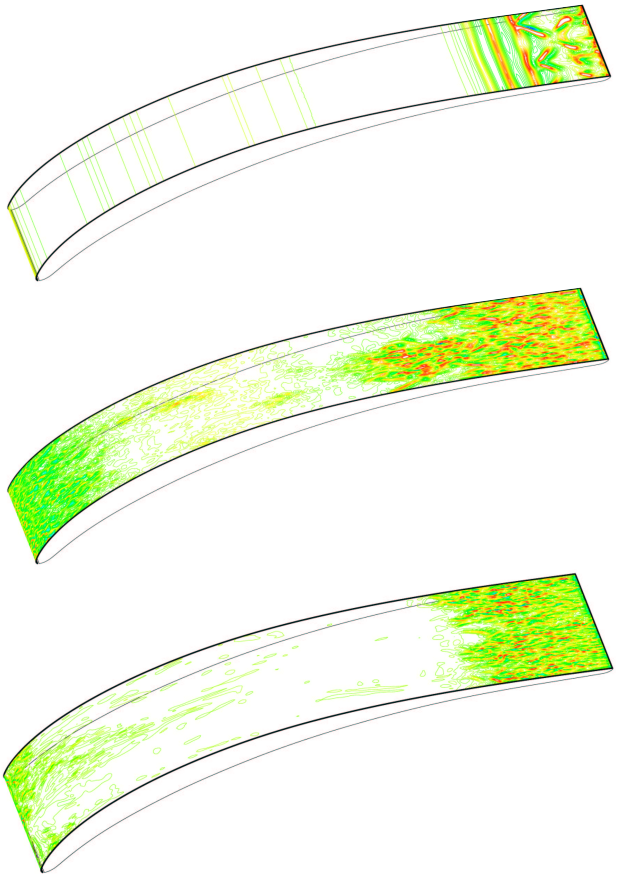

Fig. 5 Top: no inlet disturbance; middle, incident wakes; lower incident grid turbulence.

ation such that converged solutions can be obtained faster for each time step, or, alternatively, one can solve the unsteady portion of the problem in such a way that it converges faster. We are attempting to solve the inner iteration faster in two ways: firstly, we have developed a new algorithm based on an LU nonlinear Symmetric Gauss-Seidel (LU-SGS) iteration that significantly increases the convergence rate of the core solver.[1] Secondly, we have developed preconditioning approaches to improve the convergence rates of the core solver in areas where the Mach number is extremely low and the flow is nearly incompressible. We are also in the process of taking advantage of the fact that, although very complex, the nature of the unsteadiness in turbomachinery flows is periodic. Using a nonlinear frequency domain decomposition strategy, we intend to significantly speed up the convergence to a periodic steady-state.[7] Finally, we have developed a hybrid-ADI scheme for the time-advancement of the dual-time stepping scheme that speeds up the time integration and, at the same time, formally guarantees second- or third-order accuracy in time. These four algorithms have been previously described[2] and will not be presented in this paper. We will focus, however, on the development of the LU-SGS algorithm and the results and improvements that have been obtained with it so far.

LU-SGS Algorithm for the Compressible Navier-Stokes Equations

During the course of last year, new implicit algorithms have been developed for more efficient solution of the Euler equations of compressible flow, as a first step towards speed-up of Navier-Stokes solutions. The methods are based on a preconditioned, Lower-Upper (LU) implementation of a nonlinear, Symmetric Gauss-Seidel (SGS) algorithm for use as a smoothing algorithm in a multigrid method. The methods have been implemented for flows in quasi-one-dimensional ducts and for two dimensional flows past airfoils on boundary-conforming "O"-type grids for a variety of Symmetric Limited Positive (SLIP) spatial approximations, including the scalar dissipation and Convective Upwind Split Pressure (CUSP) schemes. The method appears to be significantly faster than earlier explicit or implicit methods for this class of problems, allowing solution of these problems to the level of truncation error in three to five multigrid cycles. The method has been extended to the Navier-Stokes equations, also in two dimensions, and shows significant speed-up, though not as impressive as for the Euler equations. The method also might be extended for use as a smoother in temporally sub-iterated schemes for unsteady flow problems such as the one currently used in TFLO. Further work is needed to achieve significant speedups for highly stretched Navier-Stokes meshes on complex three-dimensional configurations. This work is ongoing and has already been implemented in three dimensions for inviscid flow.

The principal ingredients of the new scheme are as follows. First, we have adopted the LU-SGS scheme in a fully nonlinear form, in which the fluxes at each cell interface are recomputed for each cell using the most recently updated values of the flow solution variables while sweeping forward from left to right, and then backward from right to left. Second, the combined Jacobian matrices that would appear on the diagonal of a linearized implicit, flux-split scheme are used as a preconditioner of the nonlinear scheme. Third, we observed a slower rate of convergence in the supersonic zones, so we introduced options for additional sweeps only over cells in which the local Mach number was supersonic. Finally, we combine this evolution algorithm with a fully nonlinear multigrid algorithm of the type formulated by Jameson,[6] which uses the full-approximation scheme (FAS) introduced by Brandt.[8]

A fast rate of convergence results from the effective combination of all these features. In the reminder of the section, results of computations of the transonic flow past several two-dimensional airfoils are

presented.

Results

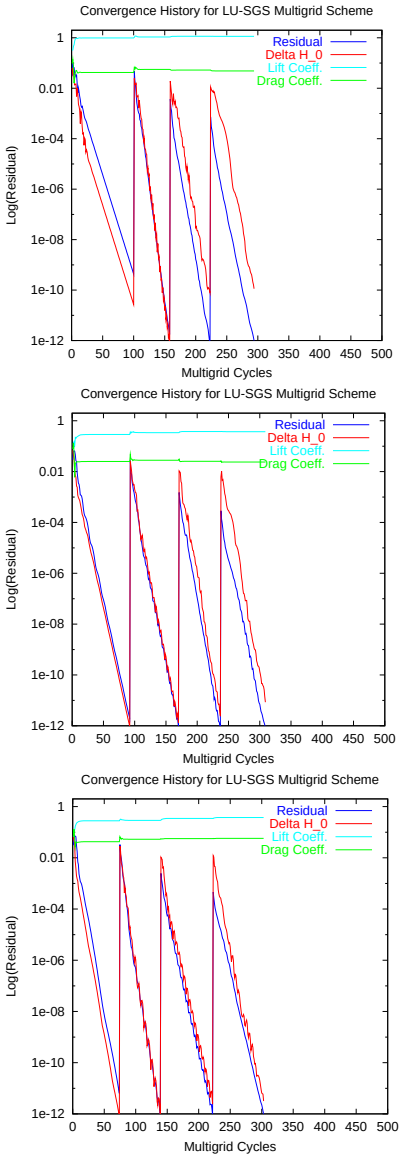

Fig. 6 Convergence rates for transonic flow past the (top) RAE 2822 airfoil at Mach 0.75 and 3.0 degrees incidence, (center) NACA 0012 airfoil at Mach 0.80 and 1.25 degrees incidence, and (bottom) NACA 0012 airfoil at Mach 0.85 and 1.00 degrees incidence. Computations proceed on a sequence of three grids, with the finest grids containing 160×32 cells.

The results of several sample computations are presented here to illustrate both, the asymptotic rate of convergence of the method, and the speed with which global measures of the solutions, such as force and moment coefficients and surface pressure distributions, converge. The airfoil flow results presented here are computed on "O"-type grids extending from the airfoil to a far-field boundary located approximately 26 chord lengths from the body surface. Results are presented for computations using only the CUSP dissipation model; comparable results for the scalar version of the SLIP scheme are presented elsewhere.[1]

All results are computed using a sequence of four grids; having 20×4, 40×8, 80×16, and 160×32 mesh cells in the wrap-around and body-normal directions, respectively. An uniform free stream is used as the initial condition for the solution on the coarsest grid, and interpolated coarser grid solutions are used as initial conditions for the subsequent, finer, grids.

Figures 6 (a), (b), and (c) show the convergence histories for the solution of three representative transonic flow problems; the flow past the RAE 2822 airfoil at Mach 0.75 and 3.00 degrees incidence, and the flow past the NACA 0012 airfoil at Mach 0.80 and 1.25 degrees incidence and at Mach 0.85 and 1.00 degrees incidence. In each case, a sequence of up to 100 multigrid cycles on each of the four successive grids is shown; the iterations are halted if the residual or the fractional deviation in total enthalpy reaches 10^{-12}. The numerical scheme is designed to drive the total enthalpy of iteratively-converged solutions to a constant, and the fractional deviation in total enthalpy is plotted in the figures along with the residual. These figures demonstrate that the asymptotic rate of convergence of the method is nearly independent of problem size (i.e., number of grid cells), and that residuals approaching 64-bit round-off levels can usually be reached in fewer than 100 multigrid cycles.

In addition to having an impressive asymptotic rate of convergence, the method also converges the global measures of error quickly. Errors in lift and drag coefficients typically are reduced to the level of truncation error (on the order of 1 - 2% of their iteratively converged values) in 3 to 5 multigrid cycles, although more cycles may be required on the coarsest grid since it is starting from a poorer initial guess. (Of course, additional computations on the coarsest grid cost very little in CPU time.)

Tabulated values of the lift and drag coefficients for these computations are compared with their (iteratively) converged values in,[1] and demonstrate that the values are within about 1% of their iteratively-converged values after only 3 multigrid cycles, and to within a fraction of 1% after only 5 cycles.

Combustor simulations

The overall goal of this group is the simulation of the reactive flow in a combustor chamber of a jet engine to study phenomena like pollutant emissions and combustion instabilities. The combustor simulations have two major components - gas phase and sprays. The gas phase part of the project has developed a parallel, unstructured grid LES solver which has now been completely integrated with the spray module. LES was chosen for its demonstrated superiority over RANS for predicting mixing, which is central to combustion.

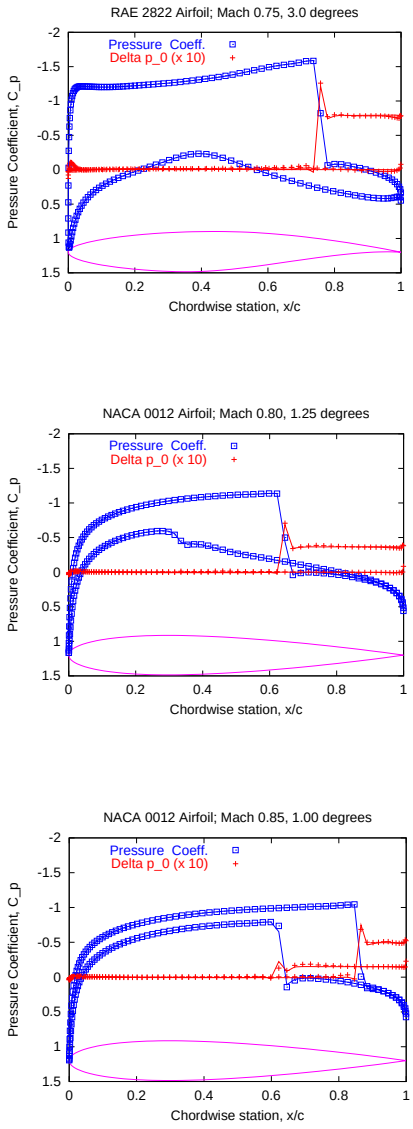

Fig. 7 Comparison of fast, preconditioned multigrid solutions after only 5 multigrid cycles on 160×32 cell grid with fully converged solutions on identical grid. (top) RAE 2822 airfoil at Mach 0.75 and 3.0 degrees incidence, (center) NACA 0012 airfoil at Mach 0.80 and 1.25 degrees incidence, (bottom) NACA 0012 airfoil at Mach 0.85 and 1.00 degrees incidence, on 160×32 cell grids.

The combustor code was christened CDP in memory of the late Dr. Charles David Pierce, who made several lasting contributions to the LES of reacting flows and specifically to this ASCI program. CDP is a code for Large Eddy Simulation of reactive flow in complex geometries. The code uses a novel approach that solves the Navier-Stokes equations in a low-Mach number form on an unstructured mesh, retaining important energy conservation properties (see Mahesh *et al.*[11,12] for details on the algorithm). A dynamic procedure[9] is used to compute the subgrid terms.

CDP Development and Validation

The code started from scratch at the beginning of the ASCI Alliance program. It has evolved during the last five years, to accommodate more general grid (it is now able to use hybrid meshes with tets, pyramids and hexs). The algorithm has also evolved during the years. It now uses a collocated mesh instead of the original staggered one and more complex combustion and spray model have been added. The time advancement is now fully implicit. The need for implicit time-advancement was felt because numerical stability restrictions imposed by the original Adams-Bashforth method were restricting the time-step in the coaxial combustor simulations to be an order of magnitude less than the time-step used by Pierce & Moin[13] in their structured grid computations that treated the viscous terms implicitly. Also the simulations performed in the front-end model showed that the narrow passages in the fuel injector considerably accelerate the flow, and thereby the convective terms impose strict restrictions on the time-step for an explicit scheme. The second-order Crank-Nicholson scheme is used for both convection and viscous terms. The convection terms are linearized prior to solution. The implementation is such that only the viscous terms can be implicitly advanced if so desired.

Results for some typical calculations are summarized in Table 1. The savings are seen to be significant.

For example, explicit calculation of the cold flow in a coaxial geometry required 320 hours $\times$ 96 processors = 30,700 CPU hours on an IBM SP3 machine. The implicit code uses around 5,000 CPU hours, which is approximately a factor of 4 times the time taken by the highly optimized structured grid solver of Pierce and Moin[13] that uses the same time-step. The structured solver is of course incapable of handling geometries as complex as the Pratt & Whitney combustor (see Fig. 8).

One of the main goals in development of the present LES code was to ensure that the solver scales well up to a high number of processors (in the range of several hundreds to thousand). Figure 9 shows the results of a scalability test on several ASCI platforms.

	Grid (10^6 cvs)	Processors	Explicit (CPU hours)	Implicit (CPU hours)
Coaxial combustor (Sommerfeld)	1.6	96	30700	5000
Turbulent channel $Re_\tau = 180$	0.9	32	2240	256
Pratt & Whitney combustor	1.4	32	13500	3200

Table 1 CPU time in hours for the the explicit and implicit solvers.

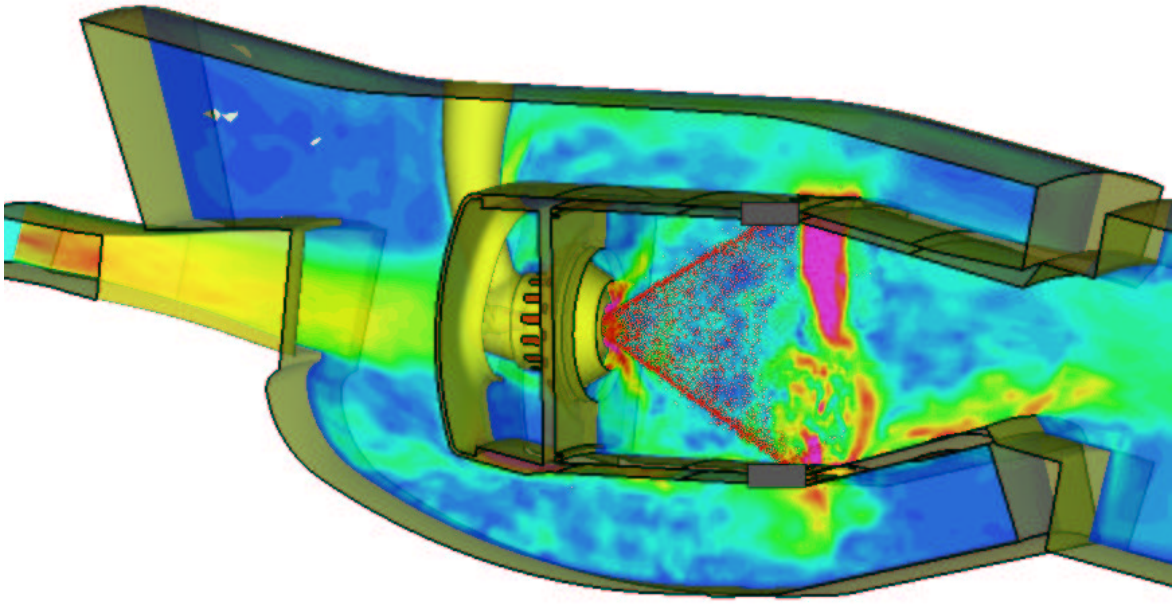

Fig. 8 Instantaneous snapshot of particles superposed on velocity contours in P&W 6000 combustor geometry

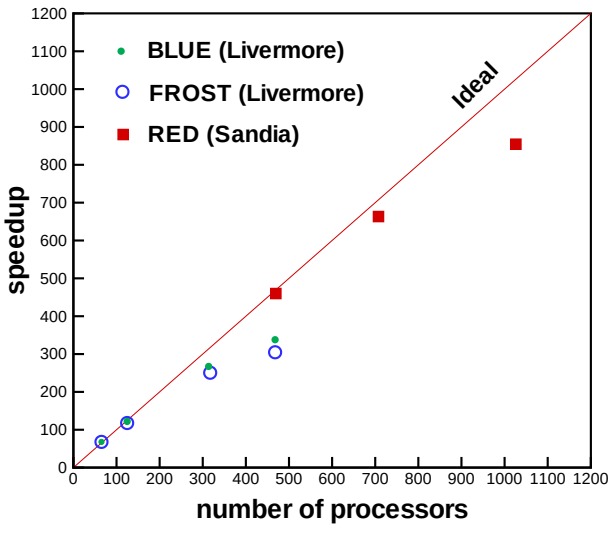

Fig. 9 Scalability test of CDP on ASCI computers for a 16 million control volume simulation

Combustion Model Implementation in CDP

CDP uses the flamelet / progress variable combustion model developed by Pierce and Moin,[14] based on "quasi-steady" flamelets in which the local flame state undergoes unsteady evolution through a sequence of stationary solutions to the flamelet equations.

A single-parameter flamelet library is first developed for the given combustor conditions by looking for stationary solutions to the one-dimensional reaction-diffusion equations. The unstable and the lower branches of the S-shaped curve are included so that the complete range of flame states, from completely extinguished (mixing without reaction) to completely reacted (equilibrium chemistry), is represented in the library. Arbitrarily complex chemical kinetic mechanisms as well as differential-diffusion effects can be included. The result is a complete set of flame states, given in terms of mixture fraction and a single flamelet parameter, denoted by λ.

In this model, in addition to the variable density, momentum, and continuity equations, scalar transport equations are solved for the mixture fraction Z, which is a conserved scalar, and for the progress variable C, which is a non-conserved scalar. The continuity equation acts as a constraint on the velocity field, with the time-derivative of density as a source term. This constraint is enforced by the pressure, in a manner analogous to the enforcement of the incompressibility constraint for constant density flows.

Under the model assumptions, all the other fluid flow variables (density, temperature, molecular viscosity, molecular diffusivity), chemical species and the reaction source terms in the scalar transport equations are related to the mixture fraction and the progress variable through a flamelet library that is precalculated given a specific fuel reaction mechanism and the flow conditions in the combustor. The only requirement for the quantity chosen to serve as the progress variable is that it should be a quantity that is representative of the overall gross flame behavior and that varies monotonically with the flame state so that its value uniquely determines it. For instance, in the reacting flow simulation of methane-air combustion in a coaxial jet combustor, the progress variable is chosen as the product mass fraction $C = y_{CO_2} + y_{H_2O}$.[14]

Spray module

The droplets are modeled as point particles which satisfy Lagrangian equations. They influence the gas phase through source terms in the gas-phase equations. As the particles move, their position is located and each particle is assigned to a control volume of the gas-phase grid. The gas-phase properties are interpolated to the particle location and the particle equations are advanced. The particles are then relocated, particles that cross interprocessor boundaries are duly transferred, source terms in the gas-phase equation are computed, and the computation is further advanced.

Stochastic Model for Secondary Breakup

Liquid spray atomization plays a crucial role in the combustion dynamics in gas-turbine combustors. In standard Lagrangian particle tracking codes, emphasis is placed on obtaining the correct spray evolution characteristics away from the injector. Only the global behavior of the primary atomization, occurring close to the injector, is considered and the details are not captured. The essential features of spray evolution, *viz.* droplet size distribution, spray angle, and penetration depth, are predicted away from the injector surface by secondary breakup models. Usually, standard deterministic breakup models based on Taylor Analogy Breakup (TAB) or wave models are employed in RANS-type computations. Liquid 'blobs' with the size of the injector diameter are introduced into the combustion chamber and undergo atomization according to the balance between aerodynamic and surface tension forces acting on the liquid phase. Both models are deterministic with 'single-scale' production of new droplets. In many combustion applications, however, injection of liquid jet takes place at high relative velocity between the two phases (high initial Weber number). Under these conditions, intriguing processes such as turbulence-induced breakup, multiple droplet collision in the dense spray region, fluctuations due to cavitating flow inside the injector, etc., contribute to the process of atomization. This results in droplet formation over a large spectrum of droplet-sizes and is not captured by the above models. Predicting the distribution of droplet sizes at each spray location is important for sheet-breakup modeling.

In order to predict the essential global features of these complex phenomena, a stochastic approach for droplet breakup which accounts for a range of product-droplet sizes has been developed by Gorokhovski & Apte.[10] Specifically, for a given control volume, the characteristic radius of droplets is assumed to be a time-dependent stochastic variable with a given initial distribution function. The breakup of parent blobs into secondary droplets is viewed as the temporal and spatial evolution of this distribution function around the parent-droplet size. This distribution function follows a certain long-time behavior, which is characterized by the dominant mechanism of breakup. The size of new droplets is then sampled from the distribution function evaluated at a typical breakup time scale of the parent drop.

Hybrid Particle-Parcel Technique for Spray Simulations

Performing spray breakup computations using Lagrangian tracking of each individual droplet gives rise to a large number of droplets ($>$ 50-100 million) very close to the injector. Computing such a large number of droplet trajectories is a formidable task even with the largest of supercomputers. In parallel computation of complex flows utilizing standard domain-decomposition techniques, the load balancing per processor is achieved by equally distributing the number of grid cells among all processors. Lagrangian particle-tracking, however, causes load-imbalance owing to the varying number of droplets per processor.

In order to overcome the above load balancing problem, the usual approach is to represent a group of droplets with similar characteristics (diameter, velocity, temperature, etc.) by a computational particle or 'parcel'. In addition, one carries the number of droplets per parcel as a parameter to be tracked. Since a parcel represents a group of droplets (of the order of 100-1000), the total number of computational particles or trajectories to be simulated is reduced significantly. With breakup, the diameter of the parcel is sampled according to the procedure given above and

the number of droplets associated with the particles is changed in order to conserve mass. This reduces the total number of computational particles per processor and increases the computational overhead with sprays by around 20% depending on the number of parcels used. Each parcel has all the droplet characteristics associated with it. The parcels methodology works well for RANS-type simulations where one is interested in time- or ensemble-averaged quantities. For LES, however, we should ideally simulate as many droplet trajectories as possible in order to obtain time-accurate results.

A hybrid scheme involving the computation of both individual droplets and parcels is proposed. The difference between droplets and parcels is simply the number of particles associated with them, N_{par}, which is unity for droplets. During injection, new particles added to the computational domain are pure drops ($N_{par} = 1$). These drops move downstream and undergo breakup according to the above breakup model and produce new droplets. This increases the number of computational particles in the domain. In the dense-spray regime, one may obtain a large number of droplets in a control volume and its immediate neighbors. The basic idea behind the hybrid-approach, is to collect all droplets in a particular control volume and to group them into bins corresponding to their size and other properties such as velocity, temperature, etc. The droplets in bins are then used to form a parcel by conserving mass, momentum and energy. The properties of the parcel are obtained by mass-weighted averaging from individual droplets in the bin. For this procedure, only those control volumes are considered for which the number of droplets increases above a certain threshold value. The number of parcels created would depend on the number of bins and the threshold value used to sample them. The parcel thus created then undergoes breakup according to the above stochastic sub-grid model. It, however, does not create new parcels. On the other hand, N_{par} is increased and the diameter is decreased by mass-conservation.

The effectiveness of this hybrid approach is demonstrated in the following computations. The implementation of this method in an unstructured LES code gives us the capability of testing and validating these models in realistic industrial geometries for various combustors with multiphase flows.

PW Frontend validation Geometry

The stochastic model along with the hybrid particle-parcel approach were used to simulate the spray evolution in the Pratt and Whitney injector. The experimental data set was obtained by mounting the injector in a cylindrical plenum through which gas with prescribed mass-flow rate was injected. The gas goes through the main and guide swirler to create a swirling jet into the atmosphere. Liquid film is injected through the filmer surface which forms an annular ring. The liquid mass-flow rate corresponds to certain operating conditions of the gas-turbine engine. Measurements of the droplet distribution and liquid mass flux in the radial direction at two different axial locations away from the injector were performed. Gas-phase statistics for mean and rms velocities are also available at these locations. The outside air-entrainment rates were measured and prescribed as inflow conditions. A snapshot of the spray evolution in the $z = 0$ plane along with the gas-phase axial velocity contours is shown in Fig. 10. The hybrid-approach used herein gives a dynamical picture with correct spray angle. Preliminary results show that the liquid mass fluxes at two downstream locations are in reasonable agreement with the experimental data. However, a longer-time sample is necessary to match the computational predictions with the experiments.

Integrated Simulations

The goal of the Multi-Component Integrated Simulation Project is to couple several components of a gas-turbine engine into a single simulation that will eventually include the entire compressor, combustor, and turbine flow paths when the computing power becomes available. The prediction of multi-component phenomena, such as compressor/combustor instability, combustor/turbine hot-streak migration, and main/secondary flow ingestion/interaction, can be improved by directly coupling the simulations of all components. Since each of the components to be coupled has to be simulated using a solver which requires parallel computing resources, it is of fundamental importance that the resulting coupled simulation be both efficient and scalable. In order to achieve this goal, we are paying close attention to the individual scalability of the component codes (TFLO and CDP) as well as to the details of the coupling and the load balancing of the resulting calculation. Current efforts have focused on coupling the gas-turbine combustor and turbine main flow path as well as the turbine main and secondary-air system flow paths.

In the turbine main and secondary-air system interaction problem, the ingestion of combustor hot-gases into the rim-cavity/secondary-flow path between the high-pressure turbine vane and blade can often lead to degradation in durability. Often the combustor hot-streak is designed to be hub-biased (higher temperatures near hub) by changing the dilution hole flows in the burner in order to avoid blade tip/BOAS (Blade Outer Air-Seal) burning. However, as a result, the hot gases from the combustor near the hub can be ingested into the gap between the high-pressure turbine vane and blade and accumulate in this rim-cavity region. This high- temperature flow can burn the vane hub trailing edge platform and even burn the high-

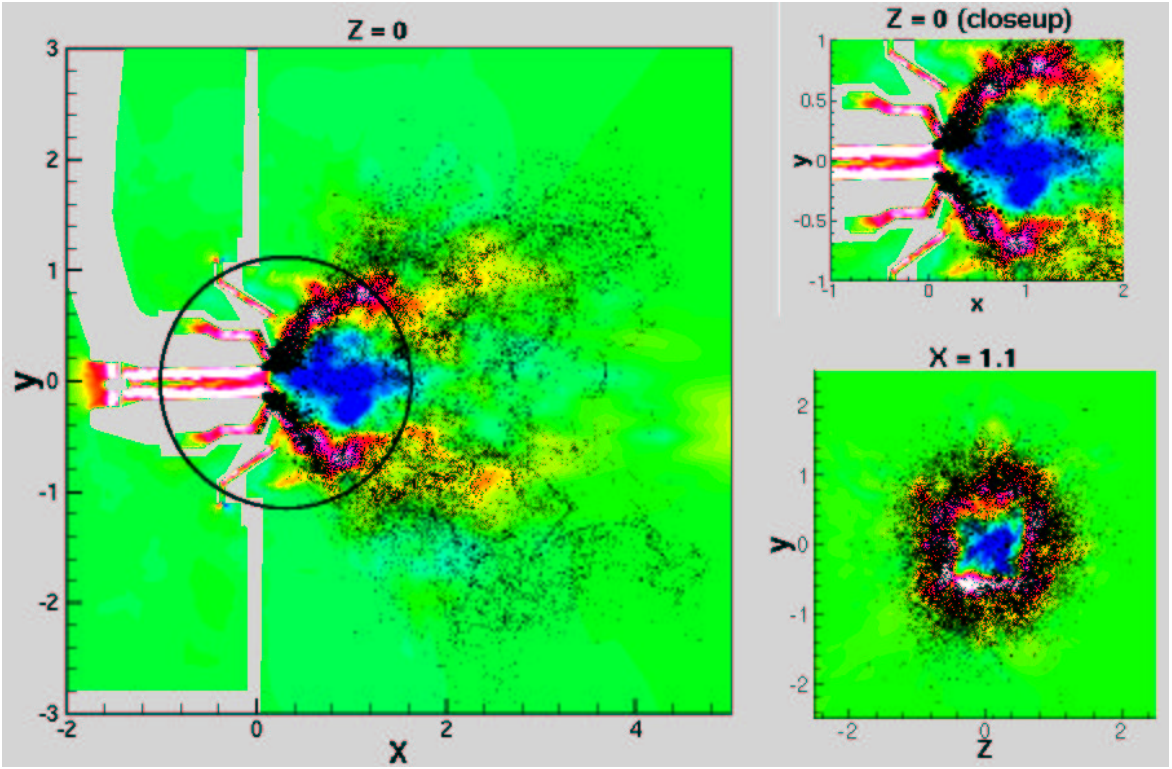

Fig. 10 Evolution of spray from a PW injector: contours of axial velocity superimposed with particle scatter plot.

pressure turbine blade disk. Disk burning can be a major concern since it can lead to premature fatigue. Accurate prediction of the flow in this region, including the amount of flow ingested or purged from the secondary-flow rim-cavity, and the temperature distributions on the vane/blade hub and blade disk, are critically important in good design.

Due to the different physics and scales in the various parts of a turbomachinery engine, it will be necessary to increase the fidelity of the physical model in areas such as the combustor. Studies have shown that capturing of the relevant physical phenomena during combustion processes can be better accomplished using Large Eddy Simulation (LES) rather than the more conventional Reynolds-Averaged Navier-Stokes (RANS) model. This increase in fidelity, however, comes at a high computational cost. From the point of view of computational simplicity it would be desirable to simulate the complete flowpath using LES. Unfortunately, this capability is beyond the computational power of the largest supercomputers in the foreseeable future. For this reason, a *zooming* approach is being followed where large portions of the flowpath that can be accurately resolved with unsteady RANS simulations (such as the compressor and the turbine) are solved with TFLO, and those portions that require LES are solved using CDP. A number of interesting research issues arise when interfacing simulations that use different physical models. Before actually coupling TFLO and CDP, we have been conducting a number of exploratory studies using TFLO and a simplified-geometry LES code to solve the fundamental issues in this coupling. Our efforts in this area are presented in this section.

Turbine Main/Secondary-Air System Integrated Simulations

The TFLO code is being used to investigate the interaction between the main flow-path and the secondary air system flow for a production HPT transonic stage, transition duct, and 1st-vane of the low-pressure turbine (LPT) which were designed based upon the rig described earlier in this paper. These simulations can be carried out using either two separate instances of the TFLO solver or a single TFLO calculation. The availability of these two options has been used to validate the multi-code interface we have developed. Thus far, both steady and unsteady integrated simulations of the main and secondary air-system flow paths have been performed. Fig. 11 shows the multi-block structured-grid used in the secondary air-system flow path. The secondary air-system computational grid is point-matched to a grid block constructed to extend across the main flow path between the end of the HPT stator hub and the beginning of the rotating HPT blade hub. Communication between the main and secondary air-system flows in TFLO does not have to be performed directly using point-matched grids. A special interface that allows point mismatched grids to be used across any block boundary[15] may be used

as an alternative to point-matched grids in order to eliminate grid stretching or skew problems that might occur. This special interface allows different computational grid systems or non-point-matched grids to communicate with each other through use of a three-dimensional interpolation process.

Figure 12 shows the temperature contours of the coupled main/secondary-air system simulation. In this simulation, the combustor hot-streak is modeled with an inlet temperature profile boundary condition to the turbine main flow-path. Since the peak temperature is near mid-span in this simulation, the temperature is high (shown in red) on the high-pressure turbine vane pressure and suction surfaces and lower on the vane hub and casing endwalls. The high-temperature flow near midspan migrates into the high-pressure turbine blade but is mixed out circumferentially due to the steady-flow inter-blade-row boundary condition treatment. The radial profile of the hot-streak is maintained, however. As a result, the temperature field in the blade passage is highest near midspan. The flow from the secondary system is at a much lower temperature than the turbine main flow path. It enters the main flow path at the hub between the high-pressure turbine vane and blade. As a result, the flow along the hub of the blade and downstream transition duct and low-pressure turbine vane is cooled. Remnants of the combustor hot-stream can be seen in Fig. 12 near the low-pressure turbine first-vane midspan position.

Flow enters the secondary air-system near the centerline and mid-way radially in the lower secondary air-system cavity. This air comes from the high-pressure compressor bleed air. In addition, air enters the tangential on-board injection (TOBI) nozzle from the combustor burner bypass. Flow exists the high-pressure secondary air-system near the centerline where it passes to the low-pressure turbine secondary air-system. In addition, flow exists between the TOBI pump and the blade disk near the top of the TOBI pump where it would proceed into the high-pressure turbine blade as film-cooling air. Finally, much of the air from the secondary air-system exits into the main flow-path at the top of the rim-cavity. This air is critical to maintain a positive purge from the secondary air-system into the main flow-path in order to keep combustor hot gases from being ingested into the rim-cavity and burning the blade disk.

RANS-LES Integrated Simulations

The integration of two flow solvers based on such different modeling approaches as the Reynolds-Averaged Navier Stokes (RANS) approach used in our turbomachinery analysis and Large Eddy Simulations (LES) used for the combustor presents a very challenging problem. In order to represent flow physics accurately, boundary conditions between zones with different models have to be carefully formulated.

Two different LES flow solvers are mainly in use at CITS. One is the unstructured CDP flow solver, which has been already described before. The other is a structured single-block flow solver written by Charles Pierce.[16] The LES formulation in this latter code is based on a low-Mach number formulation, and thus is very similar to the formulation of the CDP solver. Although a single-block structured LES flow solver is difficult to apply to the complex geometries usually found in industrial devices, the low computational cost associated with a structured LES flow solver has allowed us to explore the generation of LES boundary conditions from RANS computations in depth. Hence, it is an ideal testbed for optimizing the coupling process.

The LES routines handling the interface were written in a very generic fashion in order for them to be portable to the CDP flow solver. All communication routines follow the algorithms defined by the multi-code interface[15] which will enable compatibility between TFLO, NCC, the structured LES code and CDP.

LES Outflow Conditions from a RANS Zone

The problem of imposing RANS velocity profiles at the outflow of an LES computation has been described in.[17] The Reynolds-Averaged flow solution from the RANS computation is imposed on the developing LES solution by using virtual body forces at the outflow boundary of the LES domain:

$$F_i(\boldsymbol{x}) = \frac{1}{\tau_{\mathrm{F}}}(\bar{u}_{i,\mathrm{RANS}}(\boldsymbol{x}) - \bar{u}_{i,\mathrm{LES}}(\boldsymbol{x}))\,, \qquad (1)$$

with $\bar{u}_{i,\mathrm{RANS}}$ being the solution of the RANS flow solver computed in an overlap region between the LES and RANS domains, and $\bar{u}_{i,\mathrm{LES}}$ is the time-average of the LES solution over a trailing time-window. This body force ensures that the velocity profiles at the outlet of the LES domain fulfill the same statistical properties as the velocity profiles in an overlap region computed by a RANS simulation downstream. This makes it possible to take upstream effects of downstream flow alterations into account.

Numerical Experiment: Swirl Flow

In order to assess the accuracy and the applicability of integrated flow computations to complex flows, a swirl flow has been computed with a combination of the two flow solvers. A swirl flow was chosen, because this is typical of gas turbine combustion chambers and, due to its sensitivity to boundary conditions, it provides an excellent test-case for the accuracy of the developed boundary conditions and to validate the importance of integrated flow simulations.

In order to assess the accuracy of the LES boundary conditions, a numerical experiment was performed which is similar to the simulations performed by

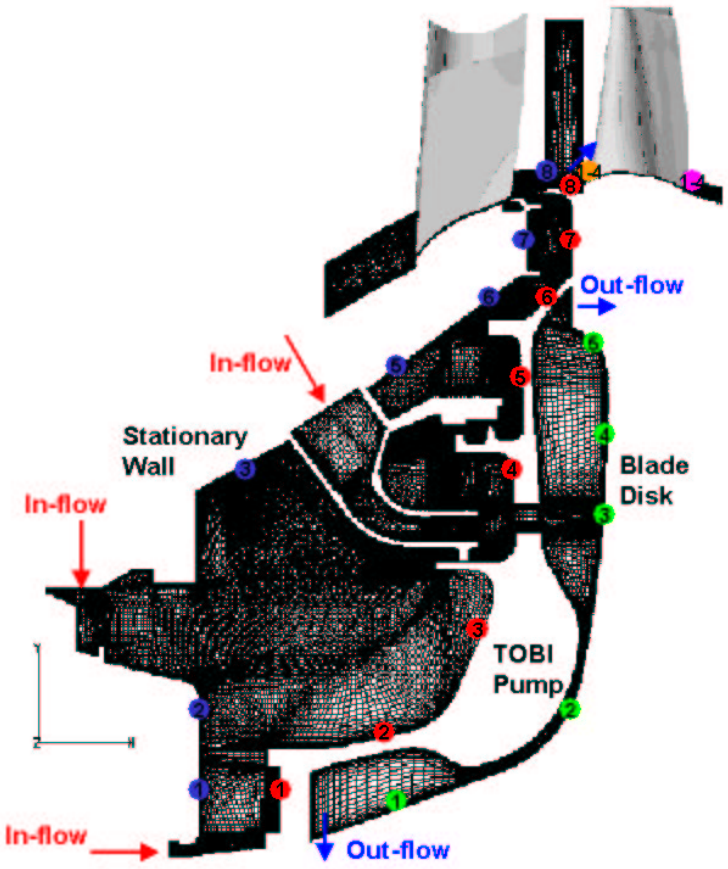

Fig. 11 Multi-Block Computational Grid Used in Integrated Main/Secondary-Air System Simulation

Schlüter *et al.*[17] to study LES outflow boundary conditions. A swirl flow at an expansion with a subsequent contraction three diameters D downstream of the expansion is considered (Fig. 13a). Inlet velocity profiles are taken from an actual experiment in a similar geometry.[18] The swirl number of the flow is $S = 0.3$, which is just supercritical, meaning that vortex breakdown takes place and a recirculation zone develops. The extension and strength of this recirculation zone is strongly influenced by the presence of the downstream contraction. The simplicity of this test-case makes it possible to perform an LES computation of the complete geometry so that an 'exact' solution of the flow field is known.

In a validation simulation, only the expansion part of the domain (Figure 13b) will be computed with the LES solver. Experience shows that the solution differs substantially if the downstream contraction is neglected.

To account for the effect of the downstream contraction, the LES flow solver is coupled with the TFLO RANS flow solver which takes care of the computation in the remaining part of the domain (Fig. 13c). If the coupling of the two codes is accurately performed, then this coupled simulation should recover the solution performed for the entire domain using the LES code alone.

Integrated LES/RANS Computations

Fig. 14 shows the velocity profiles obtained by an integrated LES-RANS computation. The swirl flow at the expansion is computed by the LES flow solver while the contraction is computed with the RANS flow solver TFLO. The velocity profiles provided by the integrated LES/RANS-computation (solid lines) are essentially identical to the velocity profiles computed by the LES of the entire domain (circles).

This computation is to the authors' knowledge the first successful integration of two separate flow solvers, one based on the RANS approach and the other on the LES approach.

Acknowledgments

This work has been supported by the Department of Energy ASCI ASAP program under contract B341491 from the Lawrence Livermore Laboratory. The authors would like to acknowledge the contributions of all the PIs at the Center for Integrated Turbulence Simulations: Profs Moin, Pitsch, Eaton, Jameson, Durbin, Davis, Mahesh, Dally and Hanrahan. Major portions of this paper represent their contributions to our 2002 Annual Report. This work has also been the product of the collaboration of many research associates, and postdoctoral and graduate students without whom the results presented in this paper would not have been

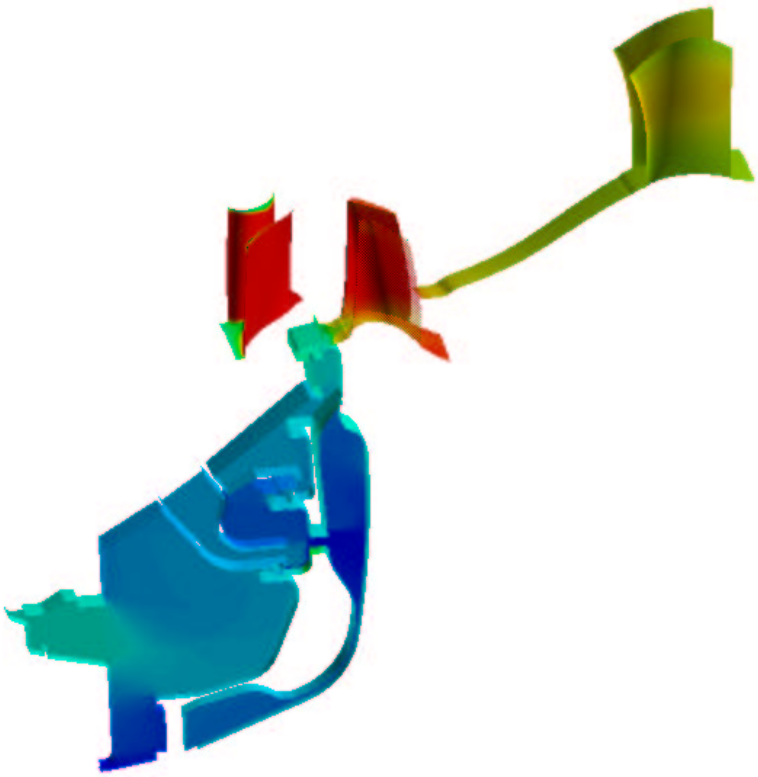

Single-Passage Sector

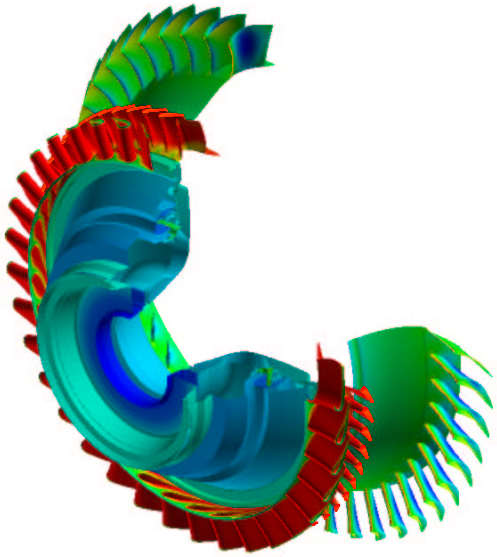

Cut-Away View of Wheel

Fig. 12 Temperature Contours of Coupled Main/Secondary-Air System Simulation

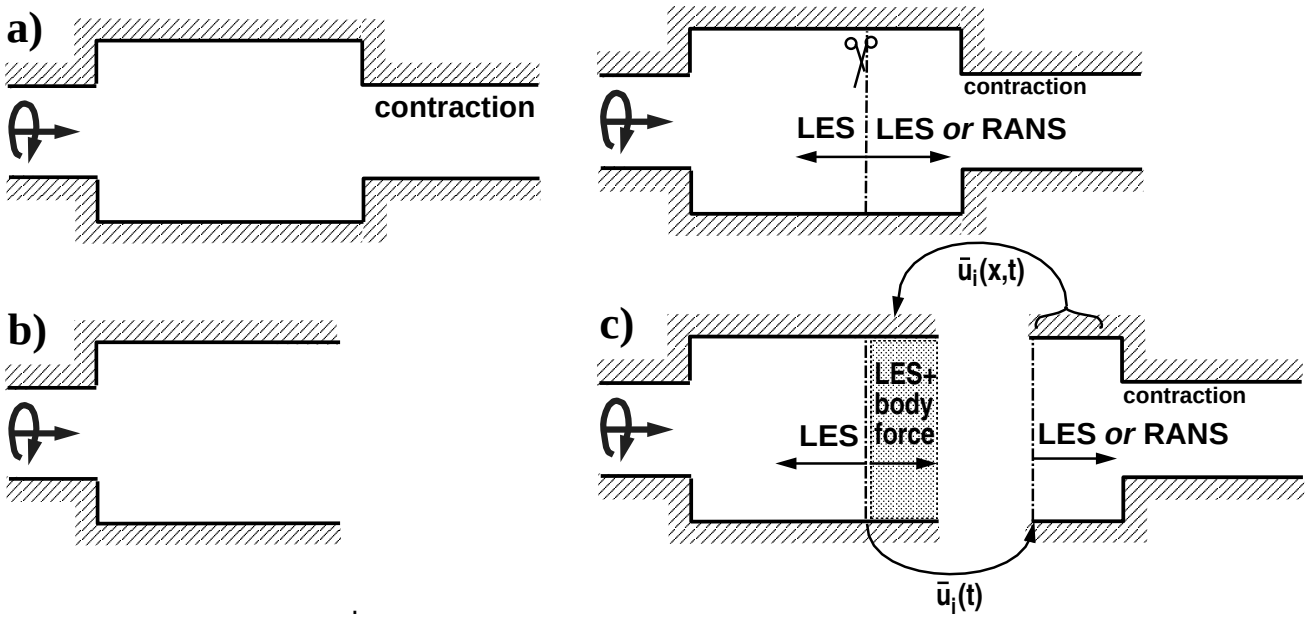

Fig. 13 Geometry for integrated LES/RANS computations: a) full geometry, b) reduced LES domain, c) schematic splitting of domain to two computational domains

possible. We recognize their efforts as fundamental to the continued success of our Center.

References

[1]Jameson, A. & Caughey, D. A., 2001, How Many Steps are Required to Solve the Euler Equations of Steady, Compressible Flow: In Search of a Fast Algorithm. *AIAA Paper 2001-2673*, AIAA 15th Computational Fluid Dynamics Conference, Anaheim, California, June 11-14, 2002.

[2]"Annual Technical Report ", 2002, Center for Integrated Turbulence Simulations, edited by W. C. Reynolds, M. Fatica and J. J. Alonso.

[3]Davis, R. L, Yao, J., Clark, J. P., Stetson, G., Alonso, J. J., Jameson, A., Haldeman, C. W., and Dunn, M. G., 2002, "Unsteady Interaction between a Transonic Turbine Stage and Downstream Components", ASME Paper GT-2002-30364.

[4]Clark, J. P., Stetson, G. M., Magge, S. S., Ni, R. H., Haldeman, C. W., and Dunn, M. G., 2000, "The Effect of Airfoil Scaling on the Predicted Unsteady Loading on the Blade of a 1 and 1/2 Stage Transonic Turbine and a Comparison with Experimental Results," ASME Paper 2000-GT-0446, ASME Turbo Expo 2000, Munich, Germany.

[5]Wu X. & Durbin P.A., 2001, Numerical simulation of heat transfer in a transitional boundary layer with passing wakes, *J. Fluid Mech.* 446:199–228.

[6]Jameson, A., 1983, Solution of the Euler Equations for Two-dimensional, Transonic Flow by a Multigrid Method. *Appl. Math. Comp.* Vol. 13, pp. 327–356.

[7]McMullen, M. S., 2003 The Application of Non-Linear Frequency Domain Methods to the Euler and Navier-Stokes Equations. Stanford University Ph.D. Thesis, March 2003.

[8]Brandt, A., 1977, Multi-Level Adaptive Solution to Boundary Value Problems, *Math. Comp.*, Vol. 31, No. 138, pp. 333-390.

[9]Germano M., Piomelli U., Moin P., and Cabot W., 1991, A dynamic subgrid-scale eddy viscosity model. *Phys. Fluids*, A(3 (7)):1760–1765.

[10]Gorokhovski, M. & Apte, S.V., 2001, Stochastic sub-grid modeling of drop breakup for LES of atomizing spray. *CTR Annual Research Briefs*, 169-176.

[11]Mahesh, K., Constantinescu G. & Moin, P., 2000, Large Eddy Simulation of gas turbine combustors. *CTR Annual Research Briefs*, 219-228.

[12]Mahesh, K., Constantinescu G., Apte, S., Iaccarino G.& Moin, P., 2001, Large Eddy Simulation of gas turbine combustors. *CTR Annual Research Briefs*, 3-17.

[13]Pierce C. D. & Moin P., 1998, Large eddy simulation of a confined jet with swirl and heat release. *AIAA Paper*, 98-2892.

[14]Pierce, C. D. & Moin, P., 2001, Progress variable approach for large eddy simulation of turbulent combustion. *Report TF - 80*, Flow Physics and Computation Division, Mechanical Engineering Dept., Stanford University, Stanford, California.

[15]Shankaran, S., Alonso, J. J., Liou M. F., Liu, N. S., Davis, R., 2001, A Multi-Code Coupling Interface for Combustor/Turbomachinery Simulations, 39th AIAA Aerospace Sciences Meeting and Exhibit, AIAA Paper 2001-0974, Reno. NV.

[16]Pierce C. D. and Moin P., 1998, Large eddy simulation of a confined coaxial jet with swirl and heat release. *AIAA Paper*, (1998-2892)

[17]Schluter J. U., Pitsch H., and Moin P., 2002, Consistent boundary conditions for integrated LES/RANS simulations: LES outflow conditions. *AIAA paper*, (2002-3121), 32nd AIAA Fluid Dynamics Conference, June 24-27 2002, St, Louis, MO.

[18]Dellenback P. A., Metzger D. E. and Neitzel G. P., 1998, Measurements in turbulent swirling flow through an abrupt axisymmetric expansion. *AIAA Journal*, 26(6):669–681.

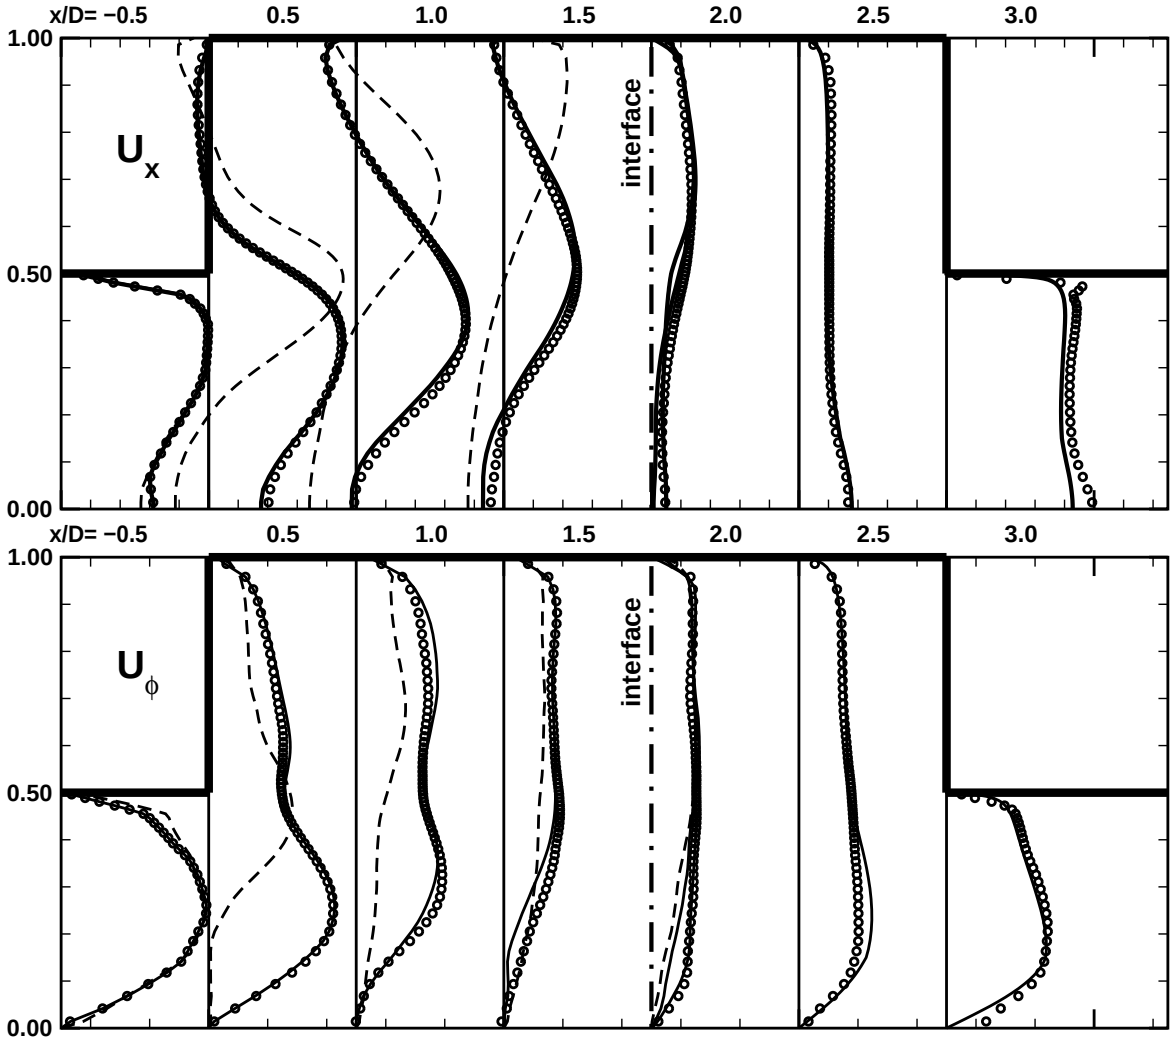

Fig. 14 Integrated LES/RANS computations. Velocity components for different downstream positions. circles: LES of full geometry (Fig. 13a), dashed line: LES of expansion (Fig. 13b), solid line: integrated LES-RANS computation (Fig. 13c)